



[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

**Search:**  The ACM Digital Library  The Guide

+write +transactions +tags +data +table +object +entry +ext

**SEARCH**

## THE ACM DIGITAL LIBRARY

[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

### Terms used

[write](#) [transactions](#) [tags](#) [data](#) [table](#) [object](#) [entry](#) [extension](#)

Found 354 of 199,986

Sort results  
by

Save results to a Binder

Search Tips

Display  
results

Open results in a new window

Try an [Advanced Search](#)

Try this search in [The ACM Guide](#)

Results 1 - 20 of 200

Result page: **1** [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale

### 1 [Cryptography and data security](#)

Dorothy Elizabeth Robling Denning  
January 1982 Book

**Publisher:** Addison-Wesley Longman Publishing Co., Inc.

Full text available: [pdf\(19.47 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

#### **From the Preface (See Front Matter for full Preface)**

Electronic computers have evolved from exiguous experimental enterprises in the 1940s to prolific practical data processing systems in the 1980s. As we have come to rely on these systems to process and store data, we have also come to wonder about their ability to protect valuable data.

Data security is the science and study of methods of protecting data in computer and communication systems from unauthorized disclosure ...

### 2 [GPGPU: general purpose computation on graphics hardware](#)

David Luebke, Mark Harris, Jens Krüger, Tim Purcell, Naga Govindaraju, Ian Buck, Cliff Woolley, Aaron Lefohn  
August 2004 **ACM SIGGRAPH 2004 Course Notes SIGGRAPH '04**

**Publisher:** ACM Press

Full text available: [pdf\(63.03 MB\)](#) Additional Information: [full citation](#), [abstract](#), [citations](#)

The graphics processor (GPU) on today's commodity video cards has evolved into an extremely powerful and flexible processor. The latest graphics architectures provide tremendous memory bandwidth and computational horsepower, with fully programmable vertex and pixel processing units that support vector operations up to full IEEE floating point precision. High level languages have emerged for graphics hardware, making this computational power accessible. Architecturally, GPUs are highly parallel s ...

### 3 [Design of the Mneme persistent object store](#)

J. Eliot B. Moss  
April 1990 **ACM Transactions on Information Systems (TOIS)**, Volume 8 Issue 2

**Publisher:** ACM Press

Full text available: [pdf\(3.22 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index](#)

[Home](#) | [Login](#) | [Logout](#) | [Access Information](#) | [Alerts](#) |

Welcome United States Patent and Trademark Office

 [Search Session History](#)[BROWSE](#)[SEARCH](#)[IEEE XPLORE GUIDE](#)

Mon, 23 Apr 2007, 5:36:06 PM EST

Edit an existing query or  
compose a new query in the  
Search Query Display.

**Search Query Display**

**Select a search number (#)**  
to:

- Add a query to the Search Query Display
- Combine search queries using AND, OR, or NOT
- Delete a search
- Run a search

**Recent Search Queries**

- #1 ( write transaction<in>metadata )
- #2 ( ( write transaction<in>metadata ) <and> ( tag<in>metadata ) )
- #3 ( ( transaction<in>metadata ) <and> ( tag<in>metadata ) )
- #4 ( ( transaction<in>metadata ) <and> ( tag<in>metadata ) )<and>  
( logical block<in>metadata )
- #5 ( ( transaction<in>metadata ) <and> ( tag<in>metadata ) )<and>  
( logical<in>metadata )
- #6 ( ( tags<in>metadata ) <and> ( transactions<in>metadata ) )  
<and> ( virtual<in>metadata )
- #7 ( ( transactions<in>metadata ) <and> ( tags<in>metadata ) )  
<and> ( identity<in>metadata )
- #8 ( ( transactions<in>metadata ) <and> ( tags<in>metadata ) )  
<and> ( identity<in>metadata )
- #9 ( write transactions<in>metadata )
- #10 ( ( data transactions<in>metadata ) <and>  
( tags<in>metadata ) )

[Help](#) [Contact Us](#) [Privacy &](#)

© Copyright 2006 IEEE -

Indexed by  
 Inspec®

terms, review

The Mneme project is an investigation of techniques for integrating programming language and database features to provide better support for cooperative, information-intensive tasks such as computer-aided software engineering. The project strategy is to implement efficient, distributed, persistent programming languages. We report here on the Mneme persistent object store, a fundamental component of the project, discussing its design and initial prototype. Mneme stores objects

**4 The design, implementation, and evaluation of Jade**

 Martin C. Rinard, Monica S. Lam

May 1998 **ACM Transactions on Programming Languages and Systems (TOPLAS)**,

Volume 20 Issue 3

Publisher: ACM Press

Full text available:  pdf(576.88 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Jade is a portable, implicitly parallel language designed for exploiting task-level concurrency. Jade programmers start with a program written in a standard serial, imperative language, then use Jade constructs to declare how parts of the program access data. The Jade implementation uses this data access information to automatically extract the concurrency and map the application onto the machine at hand. The resulting parallel execution preserves the semantics of the original serial program ...

**Keywords:** parallel computing, parallel programming languages

**5 Transactional memory: architectural support for lock-free data structures**

 Maurice Herlihy, J. Eliot B. Moss

May 1993 **ACM SIGARCH Computer Architecture News , Proceedings of the 20th annual international symposium on Computer architecture ISCA '93**, Volume 21 Issue 2

Publisher: ACM Press

Full text available:  pdf(1.10 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

A shared data structure is lock-free if its operations do not require mutual exclusion. If one process is interrupted in the middle of an operation, other processes will not be prevented from operating on that object. In highly concurrent systems, lock-free data structures avoid common problems associated with conventional locking techniques, including priority inversion, convoying, and difficulty of avoiding deadlock. This paper introduces transactional memory

**6 Minos: Architectural support for protecting control data**

 Jедидiah R. Crandall, S. Felix Wu, Frederic T. Chong

December 2006 **ACM Transactions on Architecture and Code Optimization (TACO)**,

Volume 3 Issue 4

Publisher: ACM Press

Full text available:  pdf(531.41 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We present Minos, a microarchitecture that implements Biba's low water-mark integrity policy on individual words of data. Minos stops attacks that corrupt control data to hijack program control flow, but is orthogonal to the memory model. Control data is any data that is loaded into the program counter on control-flow transfer, or any data used to calculate such data. The key is that Minos tracks the integrity of all data, but protects control flow by checking this integrity when a program uses ...

**Keywords:** Control data, buffer overflow, worms

## Secure and selective dissemination of XML documents

Elisa Bertino, Elena Ferrari

August 2002 **ACM Transactions on Information and System Security (TISSEC)**, Volume 5  
Issue 3

Publisher: ACM Press

Full text available: [pdf\(678.34 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

XML (*eXtensible Markup Language*) has emerged as a prevalent standard for document representation and exchange on the Web. It is often the case that XML documents contain information of different sensitivity degrees that must be selectively shared by (possibly large) user communities. There is thus the need for models and mechanisms enabling the specification and enforcement of access control policies for XML documents. Mechanisms are also required enabling a secure and selective dissemina ...

**Keywords:** Access control, XML, secure distribution

## 8 Logged virtual memory

Elisa Bertino, K. J. Duda

December 1995 **ACM SIGOPS Operating Systems Review , Proceedings of the fifteenth ACM symposium on Operating systems principles SOSP '95**, Volume 29  
Issue 5

Publisher: ACM Press

Full text available: [pdf\(1.52 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

## 9 Real-time shading

Marc Olano, Kurt Akeley, John C. Hart, Wolfgang Heidrich, Michael McCool, Jason L. Mitchell, Randi Rost

August 2004 **ACM SIGGRAPH 2004 Course Notes SIGGRAPH '04**

Publisher: ACM Press

Full text available: [pdf\(7.39 MB\)](#) Additional Information: [full citation](#), [abstract](#)

Real-time procedural shading was once seen as a distant dream. When the first version of this course was offered four years ago, real-time shading was possible, but only with one-of-a-kind hardware or by combining the effects of tens to hundreds of rendering passes. Today, almost every new computer comes with graphics hardware capable of interactively executing shaders of thousands to tens of thousands of instructions. This course has been redesigned to address today's real-time shading capabili ...

## 10 A structural view of the Cedar programming environment

Daniel C. Swinehart, Polle T. Zellweger, Richard J. Beach, Robert B. Hagmann

August 1986 **ACM Transactions on Programming Languages and Systems (TOPLAS)**,  
Volume 8 Issue 4

Publisher: ACM Press

Full text available: [pdf\(6.32 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper presents an overview of the Cedar programming environment, focusing on its overall structure—that is, the major components of Cedar and the way they are organized. Cedar supports the development of programs written in a single programming language, also called Cedar. Its primary purpose is to increase the productivity of programmers whose activities include experimental programming and the development of prototype software systems for a high-performance personal computer. T ...

## 11 Anatomy of a native XML base management system

T. Fiebig, S. Helmer, C.-C. Kanne, G. Moerkotte, J. Neumann, R. Schiele, T. Westmann  
 December 2002 **The VLDB Journal — The International Journal on Very Large Data**

**Bases**, Volume 11 Issue 4

**Publisher:** Springer-Verlag New York, Inc.

Full text available:  pdf(300.97 KB) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

Several alternatives to manage large XML document collections exist, ranging from file systems over relational or other database systems to specifically tailored XML base management systems. In this paper we give a tour of Natix, a database management system designed from scratch for storing and processing XML data. Contrary to the common belief that management of XML data is just another application for traditional databases like relational systems, we illustrate how almost every component in a ...

**Keywords:** Database, XML

**12 "Topologies"—distributed objects on multicomputers**

 Karsten Schwan, Win Bo

May 1990 **ACM Transactions on Computer Systems (TOCS)**, Volume 8 Issue 2

**Publisher:** ACM Press

Full text available:  pdf(3.83 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Application programs written for large-scale multicomputers with interconnection structures known to the programmer (e.g., hypercubes or meshes) use complex communication structures for connecting the applications' parallel tasks. Such structures implement a wide variety of functions, including the exchange of data or control information relevant to the task computations and/or the communications required for task synchronization, message forwarding/filtering under program control, and so o ...

**13 Parallel execution of prolog programs: a survey**

 Gopal Gupta, Enrico Pontelli, Khayri A.M. Ali, Mats Carlsson, Manuel V. Hermenegildo

July 2001 **ACM Transactions on Programming Languages and Systems (TOPLAS)**,

Volume 23 Issue 4

**Publisher:** ACM Press

Full text available:  pdf(1.95 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Since the early days of logic programming, researchers in the field realized the potential for exploitation of parallelism present in the execution of logic programs. Their high-level nature, the presence of nondeterminism, and their referential transparency, among other characteristics, make logic programs interesting candidates for obtaining speedups through parallel execution. At the same time, the fact that the typical applications of logic programming frequently involve irregular computatio ...

**Keywords:** Automatic parallelization, constraint programming, logic programming, parallelism, prolog

**14 Security: SECA: security-enhanced communication architecture**

 Joel Coburn, Srivaths Ravi, Anand Raghunathan, Srimat Chakradhar

September 2005 **Proceedings of the 2005 international conference on Compilers, architectures and synthesis for embedded systems CASES '05**

**Publisher:** ACM Press

Full text available:  pdf(396.53 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

In this work, we propose and investigate the idea of enhancing a System-on-Chip (SoC) communication architecture (the fabric that integrates system components and carries the

communication traffic between them) to facilitate higher security. We observe that a wide range of common security attacks are manifested as abnormalities in the system-level communication traffic. Therefore, the communication architecture, with its global system-level visibility, can be used to detect them. The communicati ...

**Keywords:** AMBA Bus, access control, architecture, attacks, bus, communication, digital rights management (DRM), intrusion detection, security, security-aware design, small embedded systems, system-on-chip (SoC)

**15 PELLPACK: a problem-solving environment for PDE-based applications on**

 **multicomputer platforms**

E. N. Houstis, J. R. Rice, S. Weerawarana, A. C. Catlin, P. Papachiou, K.-Y. Wang, M. Gaitatzes

March 1998 **ACM Transactions on Mathematical Software (TOMS)**, Volume 24 Issue 1

**Publisher:** ACM Press

Full text available:  [pdf\(26.30 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

The article presents the software architecture and implementation of the problem-solving environment (PSE) PELLPACK for modeling physical objects described by partial differential equations (PDEs). The scope of this PSE is broad, as PELLPACK incorporates many PDE solving systems, and some of these, in turn, include several specific PDE solving methods. Its coverage for 1D, 2D, and 3D elliptic or parabolic problems is quite broad, and it handles some hyperbolic problems. Since a PSE should p ...

**Keywords:** PDE language, execution models, knowledge bases, libraries, parallel reuse methodologies, problem-solving environments, programming-in-the-large, sofeware bus

**16 Compiler and runtime support for efficient software transactional memory**

 Ali-Reza Adl-Tabatabai, Brian T. Lewis, Vijay Menon, Brian R. Murphy, Bratin Saha, Tatiana Shpeisman

June 2006 **ACM SIGPLAN Notices , Proceedings of the 2006 ACM SIGPLAN conference on Programming language design and implementation PLDI '06**, Volume 41 Issue 6

**Publisher:** ACM Press

Full text available:  [pdf\(211.55 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Programmers have traditionally used locks to synchronize concurrent access to shared data. Lock-based synchronization, however, has well-known pitfalls: using locks for fine-grain synchronization and composing code that already uses locks are both difficult and prone to deadlock. Transactional memory provides an alternate concurrency control mechanism that avoids these pitfalls and significantly eases concurrent programming. Transactional memory language constructs have recently been proposed as ...

**Keywords:** code generation, compiler optimizations, locking, synchronization, transactional memory, virtual machines

**17 Model-driven development of Web applications: the AutoWeb system**

 Piero Fraternali, Paolo Paolini

October 2000 **ACM Transactions on Information Systems (TOIS)**, Volume 18 Issue 4

**Publisher:** ACM Press

Full text available:  [pdf\(6.94 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper describes a methodology for the development of WWW applications and a tool environment specifically tailored for the methodology. The methodology and the development environment are based upon models and techniques already used in the hypermedia, information systems, and software engineering fields, adapted and blended in an original mix. The foundation of the proposal is the conceptual design of WWW applications, using HDM-lite, a notation for the specification of structure, nav ...

**Keywords:** HTML, WWW, application, development, intranet, modeling

**18 Tradeoffs in transactional memory virtualization**

 JaeWoong Chung, Chi Cao Minh, Austen McDonald, Travis Skare, Hassan Chafi, Brian D. Carlstrom, Christos Kozyrakis, Kunle Olukotun  
October 2006 **ACM SIGPLAN Notices , ACM SIGOPS Operating Systems Review , ACM SIGARCH Computer Architecture News , Proceedings of the 12th international conference on Architectural support for programming languages and operating systems ASPLOS-XII**, Volume 41 , 40 , 34 Issue 11 , 5 , 5

**Publisher:** ACM Press

Full text available:  pdf(325.77 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

For transactional memory (TM) to achieve widespread acceptance, transactions should not be limited to the physical resources of any specific hardware implementation. TM systems should guarantee correct execution even when transactions exceed scheduling quanta, overflow the capacity of hardware caches and physical memory, or include more independent nesting levels than what is supported in hardware. Existing proposals for TM virtualization are either incomplete or rely on complex hardware impleme ...

**Keywords:** OS support, chip multi-processor, transactional memory, virtualization

**19 Model driven security: From UML models to access control infrastructures**

 David Basin, Jürgen Doser, Torsten Lodderstedt  
January 2006 **ACM Transactions on Software Engineering and Methodology (TOSEM)**, Volume 15 Issue 1

**Publisher:** ACM Press

Full text available:  pdf(968.83 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We present a new approach to building secure systems. In our approach, which we call Model Driven Security, designers specify system models along with their security requirements and use tools to automatically generate system architectures from the models, including complete, configured access control infrastructures. Rather than fixing one particular modeling language for this process, we propose a general schema for constructing such languages that combines languages for modeling systems with ...

**Keywords:** Model Driven Architecture, Object Constraint Language, Role-Based Access Control, Unified Modeling Language, metamodeling, security engineering

**20 Proxy-based acceleration of dynamically generated content on the world wide web:**

**An approach and implementation**

 Anindya Datta, Kaushik Dutta, Helen Thomas, Debra Vandermeer, Krithi Ramamritham  
June 2004 **ACM Transactions on Database Systems (TODS)**, Volume 29 Issue 2

**Publisher:** ACM Press

Full text available:  pdf(927.23 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

As Internet traffic continues to grow and websites become increasingly complex,

performance and scalability are major issues for websites. Websites are increasingly relying on dynamic content generation applications to provide website visitors with dynamic, interactive, and personalized experiences. However, dynamic content generation comes at a cost---each request requires computation as well as communication across multiple components. To address these issues, various dynamic content caching ap ...

**Keywords:** Edge caching, caching dynamically generated content, fragment caching, implementation, proxy caching, world wide web

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2007 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)